IMAGE UNDERSTANDING FOR ROBOT NAVIGATION

Robert E. Karlsen*
U.S. Army –TARDEC, Warren, MI 48397-5000

Gary Witus
Turing Associates, Ann Arbor, MI 48103

ABSTRACT

This paper presents a method to forecast terrain trafficability from visual appearance. During training, the system identifies a set of image chips (or exemplars) that span the range of terrain appearance and measures terrain trafficability characteristics as the vehicle traverses the terrain. Each chip is assigned a vector tag representing the measured vehicle-terrain interaction properties. After training, the system uses the exemplars to segment images into regions, based on visual similarity to terrain patches observed during training, and assigns the appropriate vehicle-terrain interaction tag to them. The system will therefore allow the online forecasting of vehicle performance on upcoming terrain.

1. INTRODUCTION

Most, if not all, unmanned ground vehicles currently in use are teleoperated. Typically, the operator relies exclusively on visual input from a video camera to select the route and speed. Teleoperation is robust and effective. Vision processing for autonomous and semi-autonomous navigation has not matched the human operator's visual terrain understanding. Current approaches to autonomous/semi-autonomous navigation employ a wide gamut of sensors including 3D imaging LIDAR, ground penetrating radar, multi-spectral stereo vision, ultrasound, and other sensor modalities to detect potential obstacles and forecast trafficability. Inspired by the ability of human operators, our research is focused on methods to assess terrain trafficability directly from image appearance. We do not address obstacle detection, which is an important, but separate cognitive process.

We present an approach to automated image segmentation and terrain classification using exemplars, or small image samples, to represent the variety of terrain appearance. Each chip is assigned a set of measured vehicle-terrain interaction (VTI) parameters that describe the vehicle's performance while driving over that particular terrain, and include measures such as vehicle slip, ground resistance and terrain roughness. This process requires three main functions: segmenting the terrain into areas that are visually similar, measuring and computing appropriate measures of the vehicle-terrain interaction, and matching the VTI parameters to the correct image chip.

Exemplars are used as cluster seeds to segment the terrain. Local pieces of terrain are assigned to the exemplar to which they are most similar in appearance and inherit the VTI parameters of the exemplar. Previous work has been performed in determining meaningful and robust VTI parameters (Karlsen et al., 2004). In this paper we will utilize measures of ground resistance and roughness.

Exemplar models assume that intact stimuli are stored in memory, and that classification or recognition is determined by the degree of similarity between a stimulus and the stored exemplars. Exemplar methods admit evolution of similarity metrics, since the entire sample is stored intact in memory, and not merely a feature vector summary. Simple generalization effects explain correct classification of novel (i.e., previously unseen) instances of categories. Only the item information is used for classification decisions. Categorization relies on the comparison of a new stimulus with known exemplars of the category.

Exemplar models are the most parsimonious models of categorization in terms of the underlying associative mechanism (Chase and Heinemann, 2001). Exemplar based learning was originally proposed as a model of human learning in (Medin and Schaffer, 1978), and has since been shown to explain both human and animal visual classification performance significantly better than alternative hypotheses of feature-based and prototype-based processing (Nosofsky, 1991; Werner and Rehkamper, 2001).

Various researchers have begun to develop methods to forecast traversability based on estimates of geometrical properties inferred from non-contract sensors. (Howard et al., 2001; Howard and Seraji, 2001) developed a fuzzy-rule-based system to mimic human "high/medium/low" trafficability assessment based on measures of roughness, slope and distance between obstacles computed from stereo imagery. The system was targeted for planetary rover environments. (Manchuchi et al., 2005) used a stereo color vision system together with a single axis LADAR to classify terrestrial terrain cover

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
01 NOV 2006	**N/A**	**-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Image Understanding For Robot Navigation	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **U.S. Army TARDEC, Warren, MI 48397-5000**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002075., The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**UU**	**8**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

and detect obstacles. They noted that the color-based classification system could be made more robust by considering texture of regions and shape features of objects. (Ye and Borenstein, 2004) defined a trafficability index equal to the weighted sum of the slope and roughness estimated from line-scanning laser rangefinder data. (Langer et al., 1994) classified terrain as impassible (NoGo) if any of several properties were above a threshold: height variation, the surface normal orientation, and the presence of an elevation discontinuity (all estimated from LADAR imagery). (Sarwal et al., 2003) developed a rule-based system for terrain classification from LADAR and color camera imagery.

Appearance based approaches do not attempt to directly estimate geometrical properties and then infer traversability. Instead, they classify the terrain appearance, and then assign the associated trafficability vector measured during experience traversing similar terrain. The trafficability assessment is not restricted to computations of geometrical properties, but can also reflect micro-surface properties (e.g., friction, resistance, sinkage, etc.). In previous work, we used an exemplar-based approach to segmenting terrain into Go and NoGo regions (Karlsen and Witus, 2006).

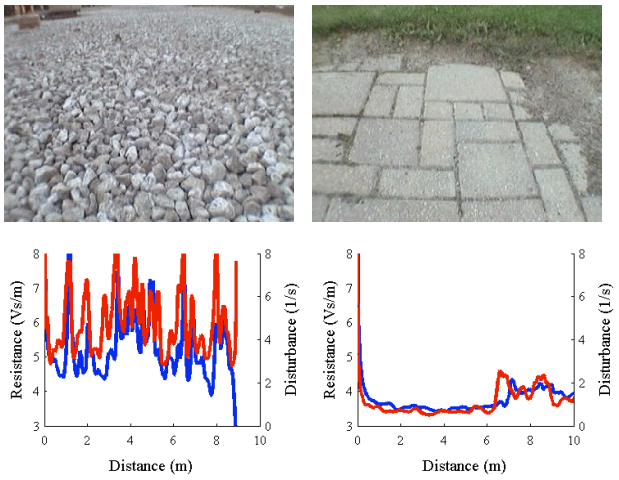

Fig. 1: Input training images and VTI parameters (resistance ($1/\alpha$) = blue, disturbance (ρ) = red).

Various applications could benefit from automatic methods to segment and classify terrain from images in addition to mobile robot navigation, such as virtual reality simulated terrain, combat engineering planning, and land cover analysis for ecological studies. These applications address different scales, terrain features and parameters of interest. It is unlikely that any specific segmentation criteria would be suitable for all of these applications. Nonetheless, the applications have important similarities. In all cases, we implicitly assume that local areas with similar appearance should be grouped together in any segmentation, and that they are likely to be representatives of the same terrain. For the purposes of this research, we assume that the segmented terrain regions do not have any a priori constraints on their geometric shape or global organization.

The approach is currently implemented as a software system designed to provide considerable flexibility in the choices of perspective transformation, resolution, scale, sampling and difference metric. In general, different choices will be appropriate for different applications.

2. TECHNICAL APPROACH

The algorithm is organized into two routines: one for offline training, which is based on fuzzy c-means clustering, and one for online learning, which applies segmentation and parameter identification to test images. At the end of the offline training, an exemplar bank is created that contains image and parameter identification data. During online learning, the exemplar bank is updated.

2.1 Training Images and Data

The user must provide a set of representative training images and associated vehicle-terrain interaction (VTI) parameters. Ideally, the training images would be drawn from the same distribution as the downstream application images. In practice, it may not be possible to ensure this. The effect on segmentation and parameter identification performance of different terrain, foliage, season, lighting, and weather between the training image set and test/application image set is a question for empirical investigation. In principle, the images can be multi-spectral with an arbitrary number of planes.

For each training sequence, a corresponding VTI data set is required that contains sensor data for the relevant patches in the imagery. If one does not have range information, assumptions, such as “flat earth,” must be made concerning the terrain in order to associate the sensor data from the vehicle to the image data that the vehicle has not traversed yet. In this paper, the vehicle is assumed to be traveling in a straight line and we use a simple estimation for the distance that each image patch is from the vehicle. For online learning and arbitrary vehicle motion, one would need to cache image patches and use a more complex method for correlating VTI parameters with image patch location. Examples of image and VTI data are shown in Figure 1.

2.2 Perspective Transformation, Resolution, Scale and Sampling

In some cases, a transformation from original camera perspective may be appropriate. In the camera image view, pixels represent the same angle (assuming lens

distortion effects are minimal), but do not project onto equal areas of ground. This is problematic since terrain appearance changes with range and thus, would require multiple instances of the same terrain for training (at different ranges).

Assuming the elevation of the camera is large relative to the variation in ground elevation in the scene, the pseudo plan view projection can be used to create a new image in which each pixel corresponds to the same ground area (see Fig. 2). The pseudo plan view projection is good for areas where the variation in elevation is small relative to the elevation of the camera, but produces distortion when this is not the case. An alternative projection is to restrict analysis to horizontal sub-bands within the image. The band view does not distort vertical objects, but retains the perspective distortion of the original camera image for flat earth regions. A third alternative is to use a stereovision camera to measure range and warp the image accordingly, such that each image chip roughly corresponds to equal areas of ground.

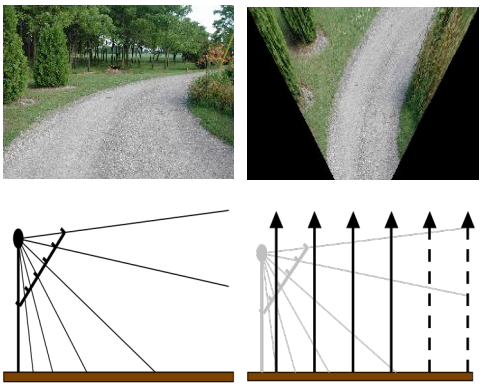

Fig 2: Camera image view and pseudo plan view.

The user must specify the analysis scale for terrain segmentation. The segmentation is based on exemplar image chips (square chips in the current software). The scale is the width of the exemplar chips. Membership in a terrain class is considered to be a bulk property of a local region, not a point-location property. The user must also specify the center-to-center spacing, or sampling distance.

2.3 Image Space Transformation

The purpose of the image space transformation is to amplify the importance of selected image properties. For example, the imagery can be transformed into a variety of color spaces. The importance of color could be strengthened or weakened by weighting different image planes. In addition to the RGB color coordinate system, we have experimented with the HSV (hue, saturation, value) and L*a*b* (luminance, red/green, yellow/blue) systems.

Another transformation option is to adjust the high spatial frequency content relative to low spatial frequency content by constructing a multi-resolution pyramid representation and then applying weights to the image planes. A common example is the Laplacian-of-Gaussian spatial bandpass pre-filtering, which is often used in stereovision processing.

The space transformation could increase the dimensionality of the image space. Consider a monocular image input. The image could be processed through a bank of N spatial filters, such as edge and corner filters at different spatial scales and orientations, with each filter producing a single-plane output image.

2.4 The Exemplar Basis Set

We use the fuzzy c-means (FCM) clustering algorithm to generate the initial exemplar basis set, since it is assumed that the system will be trained offline. In run-time operation, the online algorithm will characterize upcoming terrain, as well as generate new exemplars for unrecognized terrain.

The offline FCM algorithm processes all the images at the same time, which leads to an optimal segmentation of the images. The user chooses the number of clusters desired and the algorithm determines the cluster centers and the resulting cluster sizes. The closest image chips to the cluster centers are chosen as the exemplars and form the foundation of the exemplar bank. Various methods for determining a clustering threshold from statistics of the fuzzy clusters are being explored.

Because the online algorithm processes each image independently, it is naturally suboptimal, and therefore, additional information is considered when choosing exemplars. Here, each chip is compared to its neighbors within a specified radius to calculate the difference metric between it and each of its neighbors (the radius is a user input). The aggregate local difference between the chip and its neighbors is calculated as the weighted average of the mean and minimum differences (The weight is a user input. Weighting towards the minimum leads to a larger pool of exemplars, and weighting towards the mean leads to a smaller pool of exemplars). Chips similar to their neighbors are preferred over those that are different.

2.5 Image Chip Difference Metric

Image difference metrics remain an open issue in the evaluation of image compression schemes. While it is easy to measure the amount of compression and the encoding/decoding time, it is not clear how to measure the quality of the reconstructed image, i.e., its difference in appearance from the original. Different image characteristics are important depending on the image

content, the questions at hand, and who is looking at the image.

Similarly, there is no obviously correct metric for measuring the difference between two images. Before the images are chopped into chips, they can be processed to balance the relevant image characteristics (see 2.3 Image Space Transformation). In principle, therefore, simple measures of the aggregate difference are all that are needed. Even so, there are many different ways to calculate the difference between two image chips. Some metrics are computed from the pixel-by-pixel difference between two chips, others are calculated from the difference in statistics computed from the individual chips.

Fig 3: Reconstruction of training images from exemplars.

2.6 Output Illustration Controls

The algorithm contains options to output different images to illustrate and provide insight into the processing:

- the pseudo plan view or camera band view perspective transformation of the image;
- the exemplar chips (at their location in the image) selected from the current image;
- the segmentation of the current image based on the current bank of exemplars;
- the VTI parameters for the image; and
- the image, color coded to the VTI parameter prediction for each chip.

There is no obvious and correct way to represent the different segments for purposes of visualization. Color-coding shows the different segments, but does not give much insight into the basis for the segmentation. The software illustrates the segmentation in a way that provides direct visual insight into the basis for the segmentation. To visualize the segmentation, the software replaces each image chip with the exemplar chip to which it is associated (image chips not associated with any exemplar appear black) (See Fig. 3 for reconstruction of the images in Fig. 1). When the sampling distance is less than the exemplar scale, the exemplars are blended in the reconstruction. The visualization image is the same size as the pseudo plan view or camera band view perspective image, so it is easy to directly compare the two. By using the exemplar chips themselves, the visualization image shows what the exemplars look like, and which image chips they are associated with. Finally, comparing the visualization to the perspective image gives prima fascia evidence of the credibility of the segmentation.

3. IMPLEMENTATION

3.1 Image Processing

Image processing can be used to remove attributes of the imagery that can lead to misclassification, such as noise, color balance, and brightness. Automated features in cameras attempt to compensate for different lighting conditions and produce more life-like imagery. However, they are sometimes only partially successful, resulting in a time lag before compensation or applying the correction over the entire image when only a portion of the image needs correction. We were interested in applying a transform to the imagery such that consistent results were obtained irrespective of the lighting conditions. As an initial attempt at separating the luminance component from the color component, we tried the HSV (hue, saturation, value) color space. Although this resulted in some improvements over the RGB color space, the HSV system is unsatisfactory due to the cyclical nature of hue and the fact that HSV is far from perceptually uniform. This led to the implementation of an L*a*b* color space transform, where L^* refers to luminance and the a^* and b^* components encode the color information. The transformation to L*a*b* is nonlinear, resulting in components that are closer to perceptually uniform. All the results depicted in this paper use the L*a*b* color space transformation.

Fig 4: Images showing fluctuating lighting and automated camera response correction.

As seen in Fig. 4, our image sequences also show evidence of spurious color effects, most likely due to automated features of the camera system. The top pair was separated by approximately 1/3 second and the images on the bottom were each separated by about 1/4 second. We are considering ways to alleviate this problem, but are relying for now on the system learning these different color patterns. Based on this, we cannot rely on color to be a sufficient indicator of terrain

matching. However, one would like an unmanned vision system that is able to recognize terrain even in the presence of changing light conditions or color shifts, just like a human can.

As an initial attempt to capture other information, we have included a simple measure of texture as an additional dimension on which to differentiate and compare image exemplars. As an example, Fig. 5 shows texture planes computed with a two-dimensional standard deviation filter for the images in Fig. 1. In the future, we intend to explore whether additional texture processing or shape filtering will provide improved recognition results.

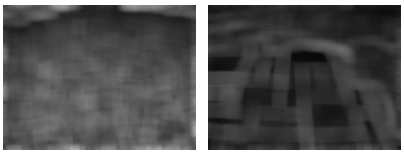

Fig 5: Texture images.

3.2 Fuzzy Clustering

For the offline portion of the system, we have implemented a fuzzy c-means (FCM) clustering algorithm (Hoppner et al., 1999). We are using the most basic form of FCM clustering with spherical clusters of the same size. Future work will look at non-spherical clusters of different sizes. As described earlier, our data consists of three L*a*b* image planes and one texture plane. The eight-element feature vector consists of the mean and standard deviation of each image chip over the four data planes. Although our system includes the ability to transform the feature vector before presenting it to the FCM algorithm, the results in this paper have no additional transform applied. In the past we have taken the square root of the data so that the FCM algorithm, which computes a root-mean-square difference, could be compared to the results of the online algorithm, which currently employs absolute differences.

The FCM algorithm provides a list of cluster centers and a matrix with the distance of each data point to each cluster. Since the cluster centers have no direct connection to the data, we move each cluster center to the location of the nearest exemplar in feature space and recompute the distances. From the resulting matrix, we can identify which exemplar (cluster center) should be assigned to each image chip in the data.

Since each chip has an associated set of VTI parameters, this gets naturally carried along with the corresponding exemplar. However, for the results in this paper, instead of using the VTI parameters for the particular exemplar, we have averaged the parameters over all chips within the cluster. There is also an option for using a weighted average, based on distance from the cluster center, where a distance of zero would yield a weight of one and a distance equal to the mean cluster distance would yield a weight of one-half. We modified the code from (Balasko et al.) for our implementation of the FCM algorithm.

3.3 Vehicle-terrain Interactions

While the test vehicle, shown in Fig. 7, had a number of sensors, we are only using a subset in this study. All VTI measures that we are interested in have a dependence on vehicle speed and therefore, this is an important parameter to measure accurately. We currently use a wheel encoder attached to a fifth wheel trailing the vehicle to provide the speed of the vehicle, irrespective of track slip. Based on previous experiments (Karlsen et al., 2004), we assume that vehicle speed is linearly proportional to the motor voltage, $v = \alpha V$. Our measure of ground resistance is the proportionality constant α (m/Vs), with high α corresponding to low ground resistance and low α to high ground resistance. Ground resistance is more directly related to $1/\alpha$, as in Fig. 1

Fig 6: Test vehicle.

The second VTI measure that we are interested in is ground disturbance. We use the output of an accelerometer positioned over the front axle to collect data and we assume that disturbance is linearly proportional to speed, $D = \rho v$. The proportionality constant ρ is used as a measure of ground disturbance, with high ρ for large disturbance and low ρ for small disturbance. The disturbance D is determined by computing the standard deviation of the accelerometer output around the point of interest.

We use a “flat earth” assumption to determine the range to points in the images. By measuring the camera height off the ground and the distance from the front axle of the vehicle to the apparent top and bottom rows of the images, we can estimate the distance to all points in the images. The vehicle was commanded to travel in a straight line. By using the distance traveled, as measured by the fifth wheel, which was synchronized with the internal vehicle sensor data, offline we can tag each image

chip with the appropriate VTI parameters. In more realistic operation, where the vehicle turns, where the terrain is not flat and where there are objects in front of the vehicle, more complex processing and data handling procedures will be required.

4. RESULTS

4.1 Data

The data collection that forms the basis for the results in this paper consists of 34 runs over different types of terrain, such as concrete, asphalt, dirt, grass, bricks, gravel, sand and rocks. Each run was between 15-25 seconds, with periods at the beginning and end where the vehicle was motionless; vehicle motion occurred for between 10-15 seconds. The vehicle-terrain interaction (VTI) parameters that we are currently exploring are for quasi-steady state conditions and so we are not considering effects due to acceleration or deceleration. For each run over a given terrain segment, we also had another run in the opposite direction.

For this paper, we chose five runs to train the system and used the companion runs in the opposite direction for testing. Terrain 1 consisted of rocks, terrain 2 was brick pavers and grass, terrain 3 was cement and grass, terrain 4 was asphalt and cement, and terrain 5 was rough sand.

We collected distance data from the fifth wheel encoders, acceleration data from the accelerometer over the front axle, voltage data from the motor, and image data from the onboard camera. The VTI parameters of interest are the ground resistance and ground disturbance.

4.2 Data Processing

We smoothed the voltage data with a Hamming-like filter of length 0.5 s. The acceleration data was converted to disturbance by a Hamming-like standard deviation filter of length 0.5 s. We used a heuristic algorithm to remove spikes from the wheel encoder data, which was then filtered by a derivative filter of length 0.5 s to produce vehicle speed. Vehicle speed divided by voltage yielded the terrain resistance parameter. Disturbance divided by vehicle speed yielded the ground disturbance parameter.

We extracted every fifth frame in the center part of each of the training sequences, resulting in 309 images. We used every frame in the center part of each of the test sequences, resulting in 1575 images. Each frame in the video was 320x240. We cropped the images to 200x160 by taking 60 pixels off each side and 80 pixels off the bottom. We chose an image chip size of 24x24, which resulted in 48 chips per frame (except for frames that included terrain that would not be traversed by the vehicle). The resulting training set had 14,808 samples and the test set had 75,560 samples. An eight-element feature vector was computed for each of the image chip samples in the training and test sets.

4.3 Training and Test Results

We chose to use 40 clusters for this test, which was based on analysis of test error results. This resulted in a training error of 7.5% and 36.1% for the ground resistance and ground disturbance predictions, respectively. The error on the test set was 11.1% and 51.5% for the ground resistance and ground disturbance, respectively. The error was computed as the absolute difference between prediction and measurement divided by the average of the two.

Fig 7: Measured vehicle-terrain interaction parameters (α = center and ρ = right).

We implemented a color-coding scheme to graphically illustrate the predicted VTI measures using the image data. The color red corresponds to the least desirable end of the parameters (0.2 for α and 2.7 for ρ), while green corresponds to the most desirable end of the parameters (0.3 for α and 0.7 for ρ). Quantities outside that range were truncated. Image chips that were determined to be too far from any exemplar were color-coded blue, with those having desirable VTI parameter values being cyan-hued and those that were least desirable were magenta-hued. The unknown chips were included in the error computations.

Figure 7 shows an example of extrapolating the measured values for the VTI parameters to specific image locations via the "flat earth" assumption, which are input to the FCM algorithm for the training image on the left. The center image contains the terrain resistance parameter and the right image contains the terrain roughness parameter. This figure illustrates one of the places where errors can enter the process: synchronizing the onboard data with the image data. Errors can enter due to faulty range estimations, but here the terrain is fairly flat and the problem is due to distances being computed from the front axle of the vehicle, whereas terrain resistance affects both the front and rear tracks, which causes the lag seen here when the vehicle is traversing a terrain boundary. This lag between terrain roughness and resistance can also be seen in the right plot of Fig. 1.

Figure 8 shows a prediction for the VTI parameters for the training image in Fig. 7 along with the color-coding scheme. The terrain resistance images are on the left and the terrain disturbance images are on the right. Note that the predictions are not very accurate, they show the pavers being rough with high resistance and the grass being smooth with low resistance. The error can be traced with the aid of the reconstruction image on the right side of Fig. 3, where a poor choice for exemplars is shown for the image chips. In this case, our training database did not contain enough samples of brick pavers in different lighting conditions. In fact, the exemplar bank contained only one exemplar from images with pavers.

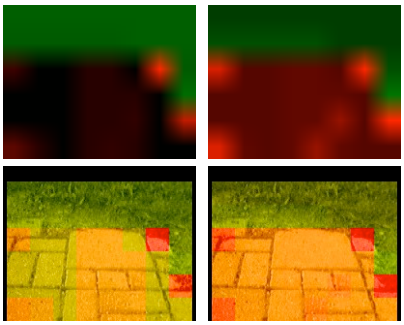

Fig 8: Predicted vehicle-terrain interaction parameters (top) and color-coded images (bottom) (α = left and ρ = right).

Figures 9-11 show further examples of test results. Fig. 9 contains data from the vehicle being run over rough rocks. Here the results are generally good and correspond to expectations for both the terrain resistance and the terrain roughness. Because of the high variability in the rock images, they tended to dominate the exemplar bank, with 16 exemplars out of 40. This is reflected in the reconstruction, which is reasonably close to the original. Exemplars derived from sand images were next with 12 exemplars.

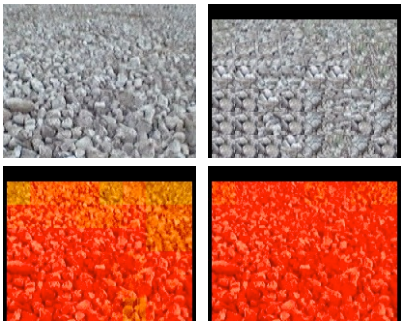

Fig. 9: Test image reconstruction (top right) and VTI predictions (bottom, α = left and ρ = right).

Figure 10 shows results for an image that contains asphalt. The system provided erroneous results for this whole image sequence since the exemplar bank did not contain any exemplars derived from asphalt, due to a shortage of asphalt images in the training sequences. As the reconstruction image shows, the system tended to pick sand exemplars as the best match, which resulted in modest agreement with the terrain resistance parameter and poor agreement for the terrain disturbance parameter.

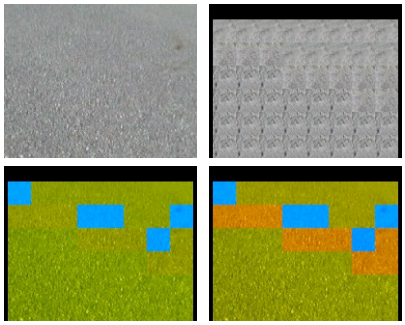

Fig. 10: Test image reconstruction (top right) and VTI predictions (bottom, α = left and ρ = right).

Figure 11 shows some interesting results for the cement images, which generally had good agreement in the cement portions, but the cracks were often mistaken for rocks, resulting in poor agreement in those specific portions.

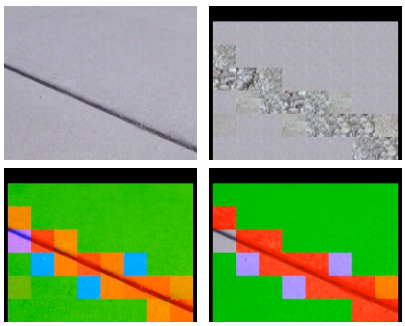

Fig. 11: Test image reconstruction (top right) and VTI predictions (bottom, α = left and ρ = right).

The pavers/grass image sequence caused the most error in the training set, while the asphalt sequence caused the most error in the test set. Creating a more balance training set should help with these errors.

5. FINDINGS AND OBSERVATIONS

This paper has demonstrated an approach to image-based terrain segmentation using exemplars, as applied to vehicle-terrain interaction (VTI) prediction. Exemplars provide a simple way to represent the characteristic color/luminance and spatial patterns of terrain. Since the exemplars are drawn from training images in such a way as to span the appearance of the training images, they are well suited to represent the variations of appearance without an a priori model of terrain appearance. The software system, as presented, allows for considerable

flexibility to specify the perspective transformation, image space transformation, scale, resolution, sampling density, and image difference metric. Empirical research is needed to tune these options for specific applications.

Preliminary results indicate the approach has potential to segment terrain in a manner that is consistent with subjective perception. The segmentation appears to provide some robustness over changes in lighting, specific terrain, and automatic camera gain and contrast adjustments, but still needs some additional work. Our previous results indicated that analysis in the camera band view was more useful for segmenting and classifying positive obstacles than the pseudo plan view. However, we have not tested the application of the pseudo plan view with this particular data set. Given that we are already using a "flat earth" assumption to extract range, the pseudo plan view may be appropriate and should be explored.

There are a number of other areas where additional work is needed. Among the most important is devising a method for determining range. The current "flat earth" assumption is not viable for real-world application. Solutions include using internal sensors to measure pitch and roll and then correct for them, although this just provides a correction for the "flat earth" assumption. More costly methods would use a stereo camera system or laser range finder. The intent of this system is not to characterize all objects in view of the vehicle and so a method to eliminate non-traversable obstacles from the imagery would be useful. This is obviously the same as an obstacle detection system and would require the same level of sophistication.

There are a number of other subjects to explore in the current system, including incorporating more complex cluster structures, although that would also require more sophisticated processing for the online algorithm. Exploring shape filtering and multi-resolution processing could yield improvements to the clustering, without much change to the existing architecture. Methods for tracking terrain segments could also prove to be valuable.

We need to explore better procedures for selecting the training data. The non-selective method currently employed results in a training set that is over-weighted in some terrain types and under-weighted in others. We also need to explore methods to further compensate for automatic gain and color distortions of the current camera system.

The system as a whole shows promise and we intend to explore further refinements in order to provide better results. The visualization tools that have been developed for this project have been very valuable in determining where the system performs correctly and where it does not and will greatly aid with upcoming enhancements.

REFERENCES

Balasko, B., J. Abonyi and B. Feil, "Fuzzy clustering and data analysis toolbox" (www.fmt.vein.hu/softcomp/fclusttoolbox).

Chase, S. and E.G. Heinemann, "Exemplar memory and discrimination," *Avian Visual Cognition*, R.G. Cook, Ed., (2001).

Hoppner, F., F. Klawonn, R. Kruse and T. Runkler, *Fuzzy Cluster Analysis: Methods for Classification, Data Analysis and Image Recognition*, John Wiley & Sons, (1999).

Howard, A., E. Tunstel, D. Edwards and A. Carlson, "Enhancing fuzzy robot navigation systems by mimicking human visual perception of natural terrain traversability," Joint 9th IFSA World Congress and 20th NAFIPS Int. Conf., Vancouver, B.C., Canada, 7-12, July 2001.

Howard, A., and H. Seraji, "Vision-based terrain characterization and traversability assessment," *J. Robotic Systems* **18**(10), 577-587 (2001).

Karlsen, R.E., J.L. Overholt and G. Witus, "Run-time assessment of vehicle-terrain interactions," Proc. 24th Army Science Conference, Orlando, FL (2004).

Karlsen, R.E. and G. Witus, "Vision-based terrain learning," *Unmanned Ground Vehicle Technology VIII*, SPIE Proc. **6230**, 33-42 (2006).

Langer, D., J.K. Rosenblatt and M. Hebert, "A behavior-based system for off-road navigation," IEEE Trans. Robotics and Automation **10**(6), 776-783 (1994).

Manduchi, R., A. Castano, A. Talukder and L. Matthies, "Obstacle detection and terrain classification for autonomous off-road navigation," *Autonomous Robots* **18**, 81-102 (2005).

Medin, D. and M. Schaffer, "Context theory of classification learning", *Psychological Review* **85**(3), 207-238 (1978).

Nosofsky, R.M., "Tests of an exemplar model for relating perceptual classification and recognition memory," *J. Experimental Psychology: Human Perceptual Performance* **17**(1), 3-27 (1991).

Sarwal, A., D. Simon and V. Rajagopalan, "Terrain classification, *Unmanned Ground Vehicle Technology V*, SPIE Proc. **5083**, 156-163 (2003).

Werner, C.W. and G. Rehkamper, "Categorization of multidimensional geometrical figures by chickens (Gallus gaullus f. domestica): fit of basic assumptions from exemplar, feature and prototype theory," *Animal Cognition* **4**, 37-48 (2001).

Ye, C. and J. Borenstein, "A method for mobile robot navigation on rough terrain," Proc. IEEE Int. Conf. Robotics and Automation, 3863-3869 (2004).